arXiv:1602.02136v1 [cs.LG] 5 Feb 2016

Reducing Runtime by Recycling Samples

Jialei Wang JIALEI@UCHICAGO.EDU
Department of Computer Science, University of Chicago, Chicago, IL 60637, USA

Hai Wang HAIWANG@TTIC.EDU
Toyota Technological Institute at Chicago, Chicago, IL 60637, USA

Nathan Srebro NATI@TTIC.EDU
Toyota Technological Institute at Chicago, Chicago, IL 60637, USA

Abstract

Contrary to the situation with stochastic gradient descent, we argue that when using stochastic methods with variance reduction, such as SDCA, SAG or SVRG, as well as their variants, it could be beneficial to reuse previously used samples instead of fresh samples, even when fresh samples are available. We demonstrate this empirically for SDCA, SAG and SVRG, studying the optimal sample size one should use, and also uncover behavior that suggests running SDCA for an integer number of epochs could be wasteful.

1. Introduction

When using a stochastic optimization approach, is it always beneficial to use all available training data, if we have enough time to do so? Is it always best to use a fresh example at each iteration, thus maximizing the number of samples used? Or is it sometimes better to revisit an old example, even if fresh examples are available?

In this paper, we revisit the notion of "more data less work" for stochastic optimization (Shalev-Shwartz & Srebro, 2008), in light of recently proposed variance-reducing stochastic optimization techniques such as SDCA (Hsieh et al., 2008; Shalev-Shwartz & Zhang, 2013), SAG (Roux et al., 2012) and SVRG (Johnson & Zhang, 2013). We consider smooth SVM-type training, i.e., regularized loss minimization for a smooth convex loss, in the data laden regime. That is, we consider a setting where we have infinite data and are limited only by time budget, and the goal is to get the best generalization (test) performance possible within the time budget (using as many examples as we would like). We then ask what is the optimal training set size to use? If we can afford making T stochastic iterations, is it always best to use $m = T$ independent training examples, or might it be beneficial to use only $m < T$ training examples, revisiting some of the examples multiple times (visiting each example T/m times on average)? Can using less training data actually improve performance (and conversely, using more data hurt performance)?

We discuss how with Stochastic Gradient Descent (SGD), there is indeed no benefit to using less data than is possible, but with variance-reducing methods such as SDCA and SAG, it might indeed be possible to gain by using a smaller training set, revisiting examples multiple times. We first present qualitative arguments focusing on the error decomposition showing this could be possible (Section 4), also revisiting the "more data less work" SGD upper bound analysis. We also conduct careful experiments with SDCA, SAG and SVRG on several standard datasets and empirically demonstrate that using a reduced training set can indeed significantly improve performance (Section 5). In analyzing these experiments, we also uncover a previously undiscovered phenomena concerning the behavior of SDCA which suggests running SDCA for an integer number of epochs could be bad, and which greatly affects the "optimal sample size" question (Section 6).

Following the presentation of SDCA, SVRG and SAG, a long list of variants and other methods with similar convergence guarantees have also been presented, including EMGD (Zhang et al., 2013), S2GD (Konečnỳ & Richtárik, 2013), Iprox-SDCA (Zhao & Zhang, 2014), Prox-SVRG (Xiao & Zhang, 2014), SAGA (Defazio et al., 2014a), Quartz (Qu et al., 2014), AccSDCA (Shalev-Shwartz & Zhang, 2014), AccProxSVRG (Nitanda, 2014), Finito (Defazio et al., 2014b), SDCA-ADMM (Suzuki, 2014), MISO (Mairal, 2015), APCG (Lin et al., 2015b), APPA (Frostig et al.,

2015a), SPDC (Zhang & Xiao, 2015), AdaSDCA (Csiba et al., 2015), Catalyst (Lin et al., 2015a), RPDG (Lan, 2015), NU-ACDM (Zhu et al., 2015), Affine SDCA and SVRG (Vainsencher et al., 2015), Batching SVRG (Babanezhad et al., 2015), and $\varepsilon\mathcal{N}$-SAGA (Hofmann et al., 2015), emphasizing the importance of these methods. We experiment with SAG, SVRG and especially SDCA as representative examples of such methods—the ideas we outline apply also to the other methods in this family.

2. Preliminaries: SVM-Type Objectives and Stochastic Optimization

Consider SVM-type training, where we learn a linear predictor by regularized empirical risk minimization with a convex loss (hinge loss for SVMs, or perhaps some other loss such as logistic or smoothed hinge). That is, learning a predictor $\mathbf{w}$ by minimizing the empirical objective:

$$\min_{\mathbf{w}} P_m(\mathbf{w}) = \frac{1}{m}\sum_{i=1}^{m} \ell(\langle \mathbf{w}, \mathbf{x}_i\rangle, y_i) + \frac{\lambda}{2}\|\mathbf{w}\|^2, \qquad (1)$$

where $\ell(z)$ is a convex surrogate loss, $\{\mathbf{x}_i, y_i\}$ are i.i.d training samples from a source (population) distribution and our goal is to get low generalization error $\mathbb{E}_{\mathbf{x},y}[err(\langle \mathbf{w}, \mathbf{x}\rangle, y)]$. Stochastic optimization, in which a single sample $\mathbf{x}_i, y_i$ (or a small mini-batch of samples) is used at each iteration, is now the dominant approach for problems of the form (1). The success of such methods has been extensively demonstrated empirically (Shalev-Shwartz et al., 2011; Hsieh et al., 2008; Bottou, 2012; Roux et al., 2012; Johnson & Zhang, 2013), and it has also been argued that stochastic optimization, and stochastic gradient descent (SGD) in particular, is in a sense optimal for the problem, when what we are concerned with is the expect generalization error (Bottou & Bousquet, 2007; Shalev-Shwartz & Srebro, 2008; Rakhlin et al., 2012; Défossez & Bach, 2015).

When using SGD to optimize (1), at each iteration we use one random training sample $(\mathbf{x}_i, y_i)$ and update $\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t - \eta g$ where $g = \nabla_{\mathbf{w}} \ell(\langle \mathbf{w}, \mathbf{x}_i\rangle, y_i) + \lambda \mathbf{w}$ is a stochastic estimation of $\nabla P_m(\mathbf{w})$ based on the single sample. In fact, we can also view g as a stochastic gradient estimation of the regularized *population* objective:

$$P(\mathbf{w}) = \mathbb{E}_{\mathbf{x},y}[\ell(\langle \mathbf{w}, \mathbf{x}\rangle, y)] + \frac{\lambda}{2}\|\mathbf{w}\|^2. \qquad (2)$$

That is, each step of SGD on the empirical objective (1), can also be viewed as an SGD step on the population objective (2). If we sample from a training set of size m without replacements, the first m iterations of SGD on the empirical objective (1), i.e., one-pass-SGD, will exactly be m iterations of SGD on the population objective (2). But, sampling with replacement from a finite set of m samples creates dependencies between the samples used in different iterations, when viewed as samples from the source (population) distribution. Such repeated use of samples harms the optimization of the population objective (2). Since the population objective better captures the expected error, it seems we would be better off using fresh samples, if we had them, rather than reusing previously used sample points, in subsequent iterations of SGD. Let us understand this observation better.

3. To Resample Or Not to Resample?

Suppose we have an infinite amount of data available. E.g., we have a way of obtaining samples on-demand very cheaply, or we have more data than we could possibly use. Instead, our limiting resource is running time. What is the best we can do with infinite data and a time-budget of T gradient calculation? One option is to run SGD on T independent and fresh samples. We can think of this as SGD on the population objective P, or as one-pass SGD (without replacement) on an empirical objective P_T (based on a training set of size T). Could it possible be better to use only $m = c \cdot T < T$ samples, for some $0 < c < 1$, and run SGD on P_m for T iterations?

3.1. SGD Likes it Fresh

One way to argue for the one-pass fresh-sample approach is that, in a worst-case sense, one-pass SGD is optimal, in that it is guaranteed to attain the best generalization error that can always be ensured (based on the norm of the data and the predictor). Using SGD with less data can only guarantee worse generalization error. Indeed, nothing we do with less data can ensure better error. However, such an argument is based on the worst-case behavior, which is rarely encountered in practice. E.g., in practice we know that multi-pass SGD (i.e., running SGD for more iterations using the same number of samples) typically do reduce the generalization error. Could we argue for fresh samples without reverting to worst-case analysis? Although doing so analytically is tricky, as our understanding of better-than-worst-case SGD behavior is very limited, we can get significant insight from considering the error decomposition.

Let us consider the effect on the generalization error of running T iterations of SGD on an empirical objective P_m based on m samples, versus T iterations of SGD on an empirical objective $P_{m'}$ based on m' $(m' < m)$ samples. The running time in both cases is the same. More importantly, the "optimization error", i.e., the sub-optimality of the empirical objective will likely be similar[1]. However, the

[1]With a smaller data set the variance of the stochastic gradient estimate is slightly reduced, but only by a factor of $1 - 1/m$,

estimation error, that is the difference between optimizing the population objective (2) and the empirical objective is lower as we have more samples. More precisely, we have that $P(\mathbf{w}_m) - \inf_{\mathbf{w}} P(\mathbf{w}) \leq O(1/(\lambda m))$ where $\mathbf{w}_m = \arg\min P_m(\mathbf{w})$ (Sridharan et al., 2009). To summarize, if using more samples, we have the same optimization error for the same runtime, but better estimation error, and can therefor expect that our predictions are better. Viewed differently, and as pointed out by (Shalev-Shwartz & Srebro, 2008), with a larger sample size we can get to the same generalization error in less time.

This indeed seems to be the case for SGD. But is it the case also for more sophisticated stochastic methods with better optimization guarantees?

3.2. Reduced Variance Stochastic Optimization

Stochastic Gradient Descent is appropriate for any objective for which we can obtain stochastic gradient estimates. E.g., we can use it directly on the expected objective (2), even if we can't actually calculate it, or its gradient, exactly. But in the past several years, several stochastic optimization methods have been introduced that are specifically designed for objectives which are finite averages, as in (1). SDCA (Hsieh et al., 2008; Shalev-Shwartz & Zhang, 2013; 2014) and SAG (Roux et al., 2012; Schmidt et al., 2013) are both stochastic optimization methods with almost identical cost-per iteration as SGD, but they maintain information on each of the m training points, in the form of dual variables or cached gradients, that help them make reduced variance steps in subsequent passes over the data (see, e.g., discussion in Johnson & Zhang, 2013, Section 4), thus improving convergence to the optimum of (1). This lead to the introduction of SVRG (Johnson & Zhang, 2013; Frostig et al., 2015b), which also reduces the variance of stochastic steps by occasionally recalculating the entire gradient (on all m training points), and achieves a similar runtime guarantee as SAG and SDCA. For both SDCA and SAG, and also for SVRG in relevant regimes. The number of iterations required to achieve a sub optimality of ϵ on (1) when the loss $\ell(\cdot)$ is smooth is

$$T = O\left(\left(\frac{1}{\lambda} + m\right)\log\frac{1}{\epsilon}\right), \tag{3}$$

compared to $O\left(\frac{1}{\lambda\epsilon}\right)$ for SGD. That is, these methods can reduce the optimization error faster than SGD, but unlike SGD their runtime depends on the sample size m. Say differently, with a smaller sample size, they can potentially obtain a smaller optimization error in the same amount of time (same number of iterations).

which might theoretically very slightly reduce the empirical optimization error. But, e.g., with over 1,000 samples the reduction is by less than a tenth of a percent and we do not believe this low order effect has any significance in practice.

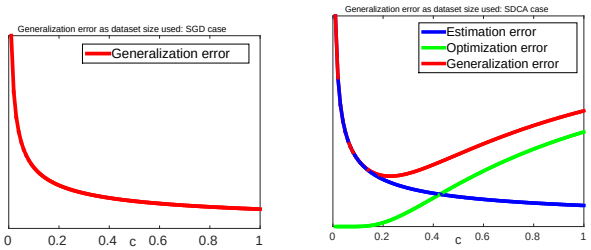

SGD *SDCA*

Figure 1. Illustration of generalization errors as c varied.

Dataset	# of instances	# of features
svmguide1	7,089	4
a9a	48,842	123
w8a	64,700	300
ijcnn1	141,691	22
covtype	581,012	54

Table 2. Statistics of datasets in this paper.

How does this affect the answer to our question? What is the best we can do with infinite data and a time-budget of T iterations with such methods? Could it be better to use only $m = c \cdot T < T$ samples, for some $0 < c < 1$?

3.3. Error Decomposition for Reduced Variance Methods

Let us revisit the error decomposition discussion from before. If we use $m' < m$ samples, the estimation error could indeed be larger. However, unlike for SGD, using less samples provides more opportunity for variance reduction, and as discussed above and can be seen from (3), can *reduce* the optimization error (or said differently, using less samples can allow us to obtain the same optimization error faster). That is, if we use $m' < m$ samples, we will have a larger estimation error, but a smaller optimization error. It might therefor be beneficial to balance these two errors, and with the right balance there is potentially for an overall decrease in the generalization error, if the decrease in the optimization error out-weights the increase in the estimation error. In Section 5 we empirically investigate the optimal sample size $m = cT$ that achieves the best balance and lowest test error, and show that it is indeed frequently beneficial to reuse examples, and that this can lead to significant reduce in test error using the same number of iterations. But first, we revisit the SGD upper bound analysis and understand what changes when we consider reduced variance methods instead.

$T\setminus$ Dataset	covtype		ijcnn1		a9a		svmguide1		w8a	
	IID	PERM	IID	PERM	IID	PERM	IID	PERM	IID	PERM
1000	0.975	0.925	0.950	0.900	1.000	0.950	0.250	0.950	1.000	0.975
2000	0.525	0.925	0.950	0.925	0.875	0.925	0.150	0.950	1.000	0.925
4000	0.375	0.950	0.650	0.950	0.825	0.925	0.125	0.975	1.000	0.975
8000	0.225	0.950	0.400	0.925	0.750	0.900	N/A	N/A	0.875	0.925
16000	0.175	0.950	0.350	0.975	0.625	0.950	N/A	N/A	0.625	0.950
32000	0.125	0.950	0.300	0.975	0.250	0.875	N/A	N/A	0.425	0.975

Table 1. The Optimal c when using SDCA under a time budget, with IID sampling and random permutation.

4. Upper Bound Analysis

In this Section, we revisit the "More Data Less Work" SGD upper bound analysis (Shalev-Shwartz & Srebro, 2008). This analysis, which is based on combining the estimation error and the SGD optimization error upper bounds, was used to argue that for SGD increasing the training set size can only reduce runtime and improve performance. We revisit the analysis considering also the optimization error upper bound (3) for the reduced variance methods. We will see that even for the reduced variance methods, relying on the norm-based upper bounds alone does not justify an improvement with a reduced sample size (i.e., a choice of $c < 1$). However, as was mentioned earlier, such an estimation error upper bound is typically too pessimistic. We will see that heuristically assuming a lower estimation error, does not justify a choice of $c < 1$ for SGD, but *does* justify it for the reduced variance methods.

The analysis is based on the existence of a "reference predictor" $\mathbf{w}_0$ with norm $\|\mathbf{w}_0\|$ and expected risk $L(\mathbf{w}_0) = \mathbb{E}[\ell(\langle \mathbf{w}_0, \mathbf{x}\rangle, y)]$ (Shalev-Shwartz & Srebro, 2008). We denote $\mathbf{w}_m$ the exact optimum of the empirical problem (1) and $\tilde{\mathbf{w}}_{\text{SGD}}$ and $\tilde{\mathbf{w}}_{\text{RV}}$ the outputs of SGD (Pegasos) and of a reduced variance stochastic method (e.g. SDCA) respectively after T iterations using a training set of size $m = cT$. The goal is to bound the generalization error of these predictors in terms of $L(\mathbf{w}_0)$, $\|\mathbf{w}_0\|$ and other explicit parameters. We assume $\|\mathbf{x}\| \leq 1$ and that the loss is 1-Lipschitz and 1-smooth.

The generalization errors can be bounded by the following error decomposition (with high probability) (Shalev-Shwartz & Srebro, 2008):

$$L(\tilde{\mathbf{w}}) - L(\mathbf{w}_0) \leq \epsilon(T) + \frac{\lambda}{2}\|\mathbf{w}_0\|^2 + O\left(\frac{1}{\lambda c T}\right) \qquad (4)$$

where $\epsilon(T) \geq P_m(\tilde{\mathbf{w}}) - P_m(\mathbf{w}_m)$ is a bound on the suboptimality of (1) (the "optimization error"), and $O(\frac{1}{\lambda c T}) \geq P(\mathbf{w}_m) - P(\mathbf{w}_0)$ is the estimation error bound(Sridharan et al., 2009). We will consider what happens when we bound the optimization error $\epsilon(T)$ as:

$$\epsilon_{\text{SGD}}(T) \leq O(1/(\lambda T))$$
$$\text{and as:} \quad \epsilon_{\text{RV}}(T) \leq \exp(-T/(1/\lambda + cT)).$$

Consider the last two terms of (4) regardless of the optimization algorithm used, even with the optimal choice $\lambda = O\left(\sqrt{\frac{1}{cT\|\mathbf{w}_0\|^2}}\right)$, these two terms are at least $O\left(\sqrt{\frac{\|\mathbf{w}_0\|^2}{cT}}\right)$, yielding an optimal choice of $c = 1$, and no improvement over one-pass SGD. This is true for both SGD and the reduced variance methods, and is not surprising, since we know that relying only on the norm of $\mathbf{w}_0$, one-pass SGD already yields the best possible guarantee—nothing will yield a better upper bound using T gradient estimates.

But the above analysis is based on a wort-case bound on the estimation error of an ℓ_2-regularized objective, which also suggests and optimal setting of $\lambda \propto 1/\sqrt{m}$ and that multiple passes of SGD (when the training set size is fixed) does not improve the generalization error over a single pass of SGD (i.e., that taking $T > m$ iterations is not any better than making $T = m$ iterations of SGD, with a fixed m). In practice, we know that the estimation error is often much lower, the optimal λ is closer to $1/m$, and that taking multiple passes of SGD certainly *does* improve performance (Shalev-Shwartz et al., 2011).

Let us consider what happens when the estimation error is small. To be concrete, let us consider a low-dimension problem where $d \ll \|\mathbf{w}_0\|^2$, though the situation would be similar if for whatever other reason the estimation error would be lower than its norm-based upper bound[2]. In d di-

[2]This could happen, for example, if low estimation error actually happens due to some other low complexity in the system, other than a bound on $\|\mathbf{w}_0\|$ and $\|\mathbf{x}\|$—either the dimensionality of the data, or perhaps the intrinsic effective dimensionality, or some combination of norm and dimensionality, or even some other norm of the data. Note that such control would have much less of an affect on the optimization, which is more tightly tied to the Euclidean norm.

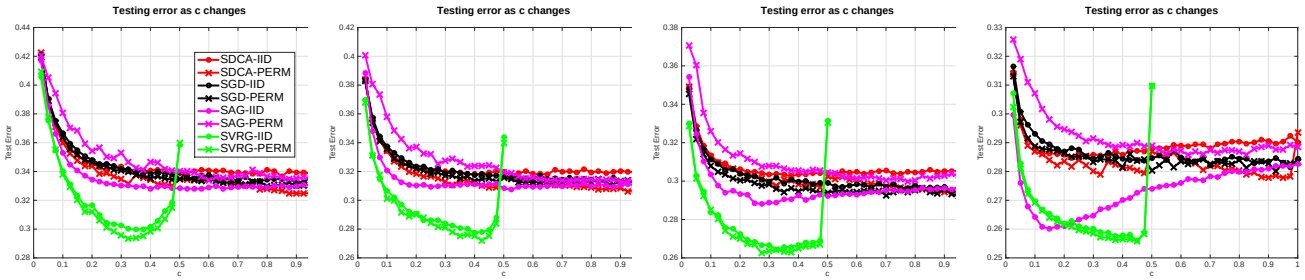

covtype, from left to right: T = 1000,2000,4000,8000

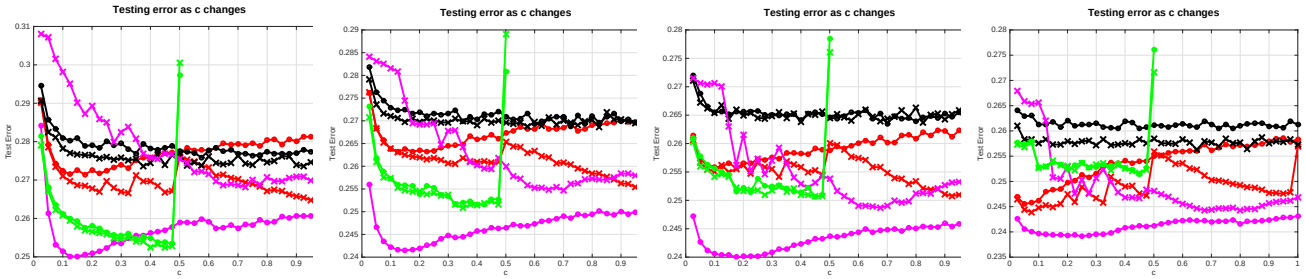

covtype, from left to right: T = 16000,32000,64000,128000

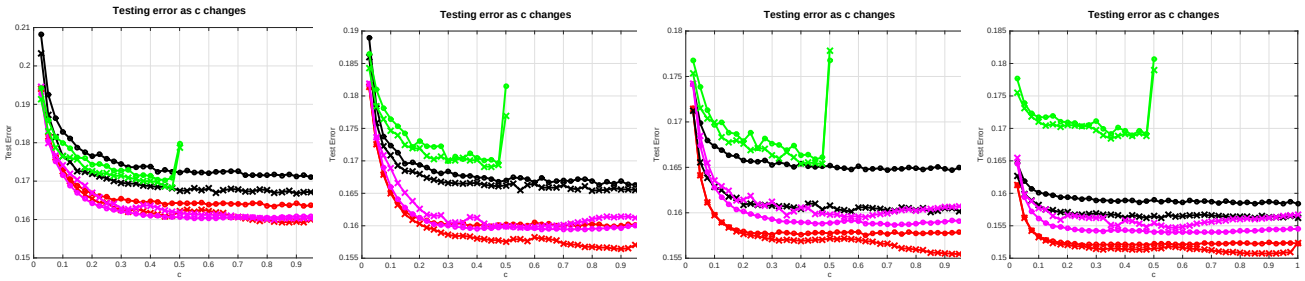

a9a, from left to right: T = 4000,8000,16000,32000

Figure 2. Illustration of generalization errors as c varied

mensions we have[3] $P_m(\mathbf{w}) - P(\mathbf{w}) \leq O\left(\sqrt{\frac{d}{m}}\right)$ yielding:

$$L(\tilde{\mathbf{w}}) - L(\mathbf{w}_0) \leq \epsilon(T) + \frac{\lambda}{2}\|\mathbf{w}_0\|^2 + O\left(\sqrt{\frac{d}{cT}}\right). \quad (5)$$

With SGD, the first two terms still yield $\Omega(\sqrt{\|\mathbf{w}_0\|/T})$ even with the best λ, and the best bound is attained for $c = 1$ (although, as observed empirically, a large range of values of c do not affect performance significantly, as the first two terms dominate the third, c-dependent term).

However, plugging in $\epsilon_{\text{RV}}(T)$, we can use a much smaller $\lambda = O\left(\sqrt{\frac{d}{cT\|\mathbf{w}_0\|^4}}\right)$ to get:

$$L(\tilde{\mathbf{w}}) - L(\mathbf{w}_0) \leq \exp\left(-T/\left(\sqrt{\frac{cT\|\mathbf{w}_0\|^4}{d}} + cT\right)\right) + O\left(\sqrt{\frac{d}{cT}}\right). \quad (6)$$

As long as $T \geq \Omega\left(\frac{\|\mathbf{w}_0\|^4}{d}\right)$, the above is optimized with $c = \Theta\left(\frac{1}{\sqrt{\log T}}\right)$ and yields:

$$L(\tilde{\mathbf{w}}) - L(\mathbf{w}_0) \leq O\left(\sqrt{\frac{d\log T}{T}}\right). \quad (7)$$

This heuristic upper bound analysis suggests that unlike for SGD, when the estimation error is smaller than its norm-based upper bound, and we are allowing a large number of iterations T, then using a reduced training set of size

[3] This is the uniform convergence guarantee of *bounded* functions with pseudo-dimension d (Pollard, 1984). Although the hinge loss is not strictly speaking bounded, what we need here is only that it is bounded at $\mathbf{w}_0$ and $\mathbf{w}$, which is not unreasonable.

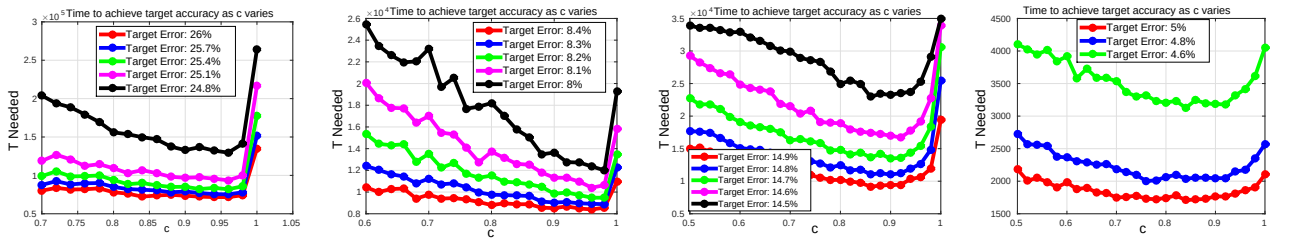

covtype *ijcnn1* *a9a* *svmguide1*

Figure 3. Illustration of the practical significance by choosing the optimal c when SDCA.

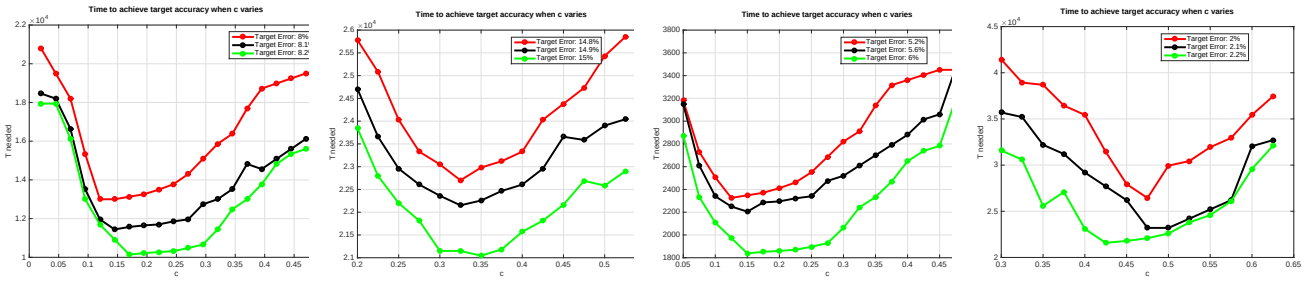

ijcnn1 *a9a* *svmguide1* *w8a*

Figure 4. Illustration of the practical significance by choosing the optimal c when SAG.

$m = cT < T$, with $c < 1$, might be beneficial. Figure 1 shows a cartoon of the error decomposition for SGD and SDCA based on this heuristic analysis.

What we have done here is revisiting the upper bound SGD analysis and understand how it might be different for reduced variance methods such as SDCA and SAG. However, this is still an upper bound analysis based on heuristic assumptions and estimation error upper bound—a precise analysis seems beyond reach using current methodology, much in the same way that we cannot quite analyze why multiple passes of SGD (for a fixed training set size) are beneficial.

5. Empirical Investigation

To investigate the benefit of using a reduced training set empirically, we conducted experiments with SDCA, SAG and SVRG (and also SGD/Peagsos) on the five datasets described in Table 2, downloaded from the LIBSVM website (Chang & Lin, 2011). We first fixed the time budget T, and randomly sampled cT instances from the data pool. Then we ran SGD, SDCA, SAG and SVRG with T iterations on the cT sample, and tested the classification performance in an unseen test dataset which consists 30% of the total instances. A value of $c = 1$ corresponds to using all fresh samples, while with $c < 1$ we reuse some samples. We tried $c = \{0.025, 0.05, 0.075, ...0.975, 1\}$, and for every setting of c and T, we follow the same protocol that optimizing λ to achieve the best prediction performance on test dataset (following Shalev-Shwartz & Srebro (2008)). For the all these algorithms, we tried both i.i.d sampling (with replacement), as well as using a (fresh) random permutation over the training set in each epoch, thus avoiding repeated samples inside an epoch. Although most theoretical guarantees are for i.i.d. sampling, Such random-permutation sampling is unknown to typically converge faster than i.i.d sampling and is often used in practice (see Recht & Ré 2012; Gürbüzbalaban et al. 2015 for recent attempts at analyzing random permutation sampling). All datasets are prepared for binary classification problem and we used the smoothed hinge loss. To overcome randomness, we repeat our experiments 500 times and report the average classification error. [4]

The results with SDCA, SAG and SVRG are shown in Figure 2 (see also additional plots in appendix), where we plot the test error as a function of the parameter c (training set size as ratio of number of iterations), while fix the time budget (number of iterations) T, and in Table 1, 3, 4, 5 where we summarize the optimal c. On all datasets, the optimal c for large enough T is less than 1. The advantage of using $c < 1$ and resampling data is more significant on covtype and svmguide1, which are both low dimensional, matching

[4] In both SAG and SVRG algorithm, a constant stepsize is used in different iterations. To obtain the best performance, we tune the stepsize for each dataset and T combination. In SVRG, one pass SGD is used to initialize, and we set $m = 2n$.

$T\backslash$ Dataset	covtype			ijcnn1			w8a		
	c	error	error (c=1)	c	error	error (c=1)	c	error	error (c=1)
1000	1.000	0.328	0.328	0.125	0.096	0.097	0.575	0.027	0.028
2000	0.525	0.307	0.311	0.100	0.092	0.094	0.55	0.026	0.026
4000	0.275	0.288	0.298	0.100	0.089	0.094	0.350	0.025	0.025
8000	0.150	0.260	0.283	0.075	0.088	0.091	0.275	0.022	0.024
16000	0.125	0.250	0.261	0.075	0.084	0.089	0.300	0.021	0.023
32000	0.150	0.242	0.251	0.025	0.082	0.083	0.225	0.018	0.020

Table 3. The Optimal c and their test error when using SAG under a time budget with IID sampling.

$T\backslash$ Dataset	covtype			ijcnn1			a9a		
	c	error	error (c=0.5)	c	error	error (c=0.5)	c	error	error (c=0.5)
1000	0.350	0.300	0.358	0.350	0.082	0.098	0.475	0.181	0.193
2000	0.400	0.278	0.344	0.400	0.072	0.091	0.475	0.178	0.188
4000	0.325	0.264	0.331	0.475	0.070	0.087	0.450	0.170	0.180
8000	0.450	0.256	0.310	0.425	0.068	0.083	0.475	0.170	0.182
16000	0.475	0.253	0.297	0.350	0.066	0.084	0.450	0.166	0.177
32000	0.425	0.252	0.281	0.375	0.066	0.083	0.425	0.169	0.175

Table 4. The Optimal c when using SVRG under a time budget with IID sampling.

the theory.

Another way of looking at the same results is asking "what is the runtime required to achieve a certain target accuracy?". For various target accuracies and each value of c, we plot in Figure 3, 4 the minimal T such that using cT samples and T iterations achieves the desired accuracy. Viewed this way, we see how using less data can indeed reduce runtime.

In SDCA, both with i.i.d and random permutation sampling we often benefit from $c < 1$. Not surprising, sampling "without replacement" (random permutation sampling) is generally better. But the behavior for random permutation sampling is particularly peculiar, with the optimal c always very close to 1, and with multi-modal behavior with modes in inverse integers, $c = 1, 1/2, 1/3, 1/4, \ldots$. To understand this better, we looked more carefully at the behavior of SDCA iterations.

6. A Closer Look at SDCA-Perm

In this section, we explore why for SDCA with random permutation, the optimal c is usually just below 1 (around $0.9 < c < 1$). We show a previously unexplained behavior of SDCA-Perm (i.e. using an independent random permutation at each epoch) that could be useful for understanding the test error as c changes. All theoretical analysis of SDCA we are aware of are of i.i.d. sampling (with replacement), and although known to work well in practice, not much is understood theoretically on SDCA-Perm. Here we show its behavior is more complex than what might be expected.

Many empirical studies of SDCA plot the sub-optimality only after integer numbers of epochs. Furthermore, often only the dual, or duality gap, is investigated. Here we study the detailed behavior of the primal suboptimality, especially at the epoch transition period. We experimented with the same datasets as used in previous section, randomly choose a $m = 4000$ subset (we observe the same experimental phenomenon for all dataset size, here we report on subsets of size $m = 4000$ for simplicity). We test with λ: 0.1 and 0.01 (the optimal regularization lies between these two values). We ran SDCA-IID and SDCA-Perm 500 times in Figure 5 plot the average behavior across the runs: the primal sub-optimality, dual-optimality and the duality gap of the iterates. We observe that:

- The behavior of SDCA-IID is as expected monotonic and mostly linear. Also, as is well know, SDCA-Perm usually converges faster than the SDCA-IID after the first epoch.
- SDCA-Perm displays a periodic behavior at each epoch with opposite behaviors for the primal and dual suboptimaitiesl: the primal decreases quickly at the beginning of the epoch, but is then flat and sometimes even increases toward the end of the epoch; The dual suboptimality usually decreases slowly at the beginning, but then drops toward the end of the epoch.

This striking phenomena is consistent across data sets. The periodic behavior explains why for SDCA-Perm the op-

T\ Dataset	covtype			ijcnn1			a9a		
	c	error	error (c=0.5)	c	error	error (c=0.5)	c	error	error (c=0.5)
1000	0.325	0.293	0.360	0.350	0.081	0.094	0.475	0.178	0.190
2000	0.425	0.272	0.340	0.425	0.072	0.091	0.450	0.176	0.186
4000	0.275	0.263	0.330	0.450	0.071	0.087	0.475	0.168	0.179
8000	0.450	0.256	0.310	0.400	0.067	0.085	0.450	0.169	0.177
16000	0.450	0.252	0.301	0.475	0.066	0.082	0.475	0.165	0.178
32000	0.350	0.251	0.289	0.350	0.066	0.082	0.350	0.168	0.172

Table 5. The Optimal c when using SVRG under a time budget with permutation.

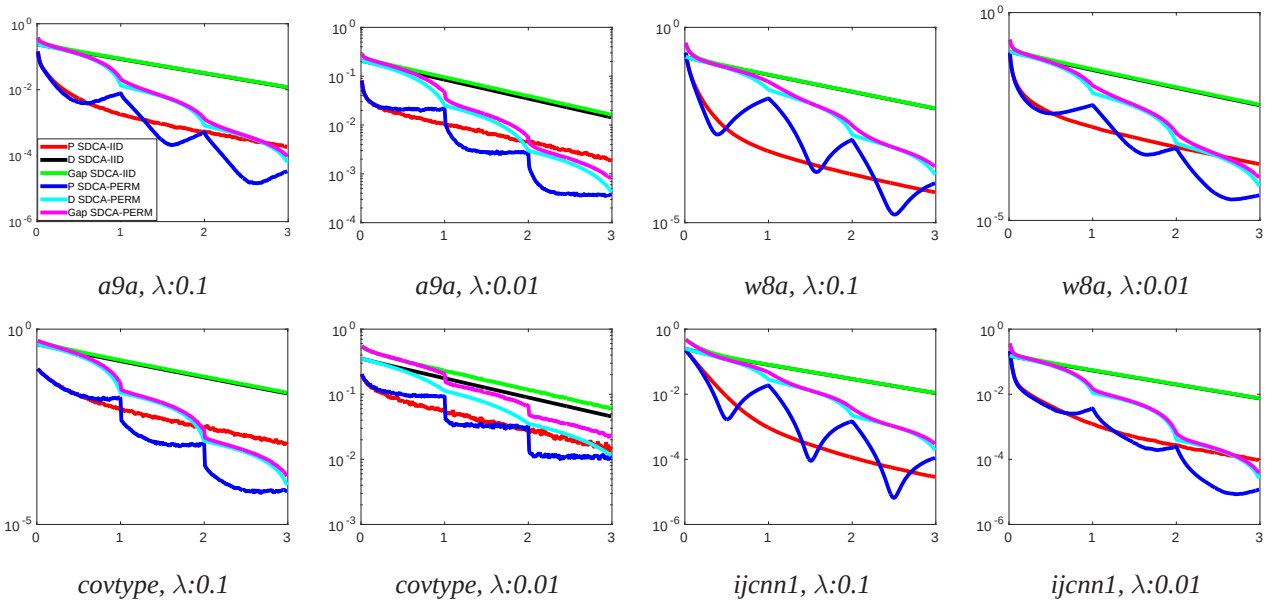

Figure 5. The convergence behavior of SDCA-Perm

timal c is usually between 0.9 and 1: since the primal improves mostly at the beginning of an epoch, we will prefer to run SDCA-Perm just more than integer number of epochs to obtain low optimization error. Returning to Figure 2, we can further see that the locally best c, for SDCA-Perm are indeed just lower than integer fractions (just before $1/2, 1/3, 1/4$ etc), again corresponding to running SDCA-Perm for a bit more than an integer number of epochs.

To understand the source of this phenomena, consider the following construction: A data set with 10 data points in $\mathbb{R}^{11}$, where each data point $\mathbf{x}_i$ has two non-zero entries: a value of 1 in coordinate i, and a random sign at the last coordinate. The corresponding label y_i is set to the last coordinate of $\mathbf{x}_i$. Let us understand the behavior of SDCA on this dataset. In Figure 6(a-b) we plot the behavior of SDCA-Perm on such synthetic data, as well the behavior of SDCA-Cyclic. SDCA-Cyclic is a deterministic (and thus easier to study) variant where we cycle through the training examples in order instead of using a different random permutation at each iterations. We can observe the phenomena for both variants, and will focus on SDCA-Cyclic for simplicity. In Figure 6(c) we plot the loss and norm parts of the primal objective separately, and observe that the increase in the primal objective at the end of each epoch is due to an increase in the norm without any reduction in the loss. To understand why this happens, we plot the values of the 10 dual variables at the end of each epoch (recall that the variables are updated in order). The first variables updates at each epoch are set to rather large values, larger than their values at the optimum, since such a value is optimal when the other dual variables are zero. However, once other variables are increased, in order to reduce the norm, the initial variables set must be decreased—this is not possible without revisiting them again. Although real data sets are not as extreme case, it seems that such a phenomena do happen also there.

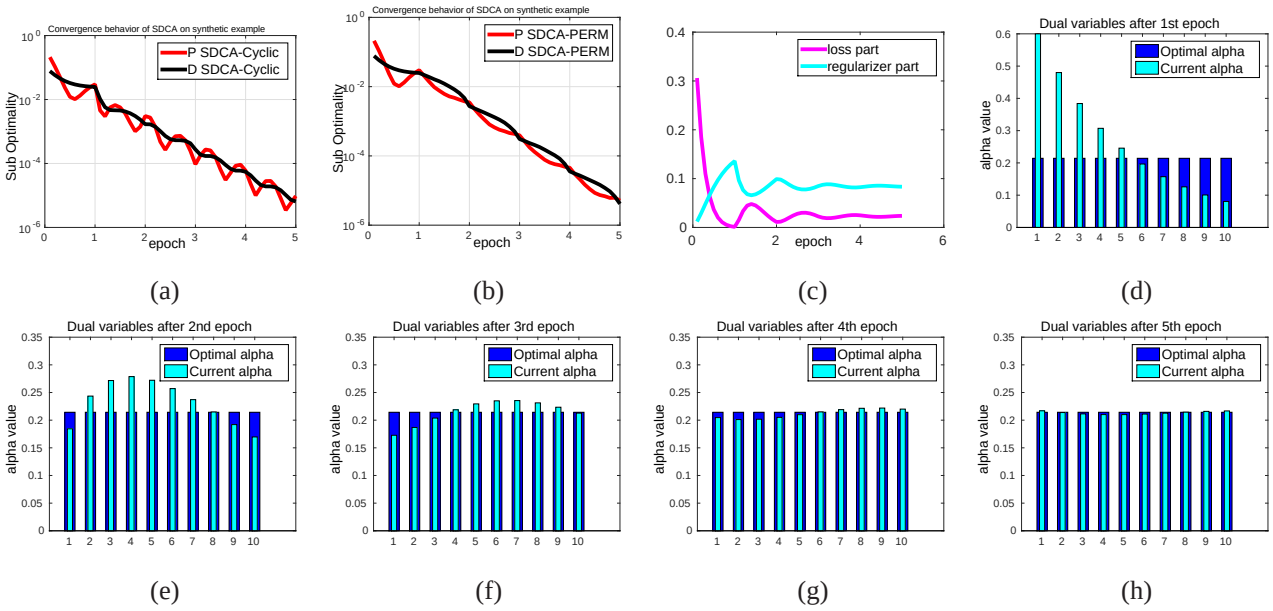

Figure 6. A synthetic example to demonstrate to behavior of SDCA

7. Conclusion

We have shown that contrary to Stochastic Gradient Descent, when using variance reducing stochastic optimization approaches, it might be beneficial to use less samples in order to make more than one pass over (some of) the training data. This behavior is qualitatively different from the observation made about SGD where using more samples can only reduce error and runtime. Furthermore, we showed that the optimal training set size (i.e., optimal amount of recycling) for SDCA with random permutation sampling (so-called "sampling without replacement") rests heavily on a previously undiscovered phenomena that we uncover here.

Our observations provide empirical guidance for using SDCA, SAG and SVRG:

First, it suggests that even when data is plentiful, it might be beneficial to use a limited training set size in order to reduce runtime or improve accuracy after a fixed number of iterations. For SDCA-Perm , it seems that the optimal strategy is often to use a slightly smaller training set than the maximal possible, and for SVRG the optimal strategy is to use a m slightly smaller than $T/2$. For SAG the optimal number of examples is more variable. Our observations are mostly empirical, backed only by qualitative reasoning—obtaining a firmer understanding with more specific guidelines of the optimal number of samples to use would be desirable.

Second, the behavior of the SDCA primal objective that we uncover suggests that performing an integer number of epochs (passes over the data), as is frequently done in practice and is the default for most SDCA packages, can significantly hurt the performance of SDCA. This is true regardless of whether we are in a data-laden regime or in a data-limited regime where we are performing multiple passes out of necessity. Instead, our observations suggest it is often advantageous to perform a few more iterations into the next epochs in order to significantly improve the solution. Further understanding of the non-monotone SDCA behavior is certainly desirable (and challenging), and we hope that pointing out the phenomena can lead to further research on understanding it, and then to devising improved methods with more sensible behavior.

References

Babanezhad, Reza, Ahmed, Mohamed Osama, Virani, Alim, Schmidt, Mark, Konečnỳ, Jakub, and Sallinen, Scott. Stop wasting my gradients: Practical svrg. *NIPS*, 2015.

Bottou, Léon. Stochastic gradient tricks. In Montavon, Grégoire, Orr, Genevieve B., and Müller, Klaus-Robert (eds.), *Neural Networks, Tricks of the Trade, Reloaded*, Lecture Notes in Computer Science (LNCS 7700), pp. 430–445. Springer, 2012.

Bottou, Léon and Bousquet, Olivier. The tradeoffs of large scale learning. In *NIPS*, pp. 161–168, 2007.

Chang, Chih-Chung and Lin, Chih-Jen. Libsvm: A library for support vector machines. *ACM Trans. Intell. Syst. Technol.*, 2(3):27:1–27:27, May 2011. ISSN 2157-6904.

Csiba, Dominik, Qu, Zheng, and Richtarik, Peter. Stochastic dual coordinate ascent with adaptive probabilities. In *ICML*, 2015.

Defazio, Aaron, Bach, Francis, and Lacoste-Julien, Simon. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *NIPS*, pp. 1646–1654, 2014a.

Defazio, Aaron, Domke, Justin, and Caetano, Tiberio. Finito: A faster, permutable incremental gradient method for big data problems. In *ICML*, pp. 1125–1133, 2014b.

Défossez, Alexandre and Bach, Francis R. Averaged least-mean-squares: Bias-variance trade-offs and optimal sampling distributions. In *AISTATS*, 2015.

Frostig, Roy, Ge, Rong, Kakade, Sham, and Sidford, Aaron. Un-regularizing: approximate proximal point and faster stochastic algorithms for empirical risk minimization. In *ICML*, pp. 2540–2548, 2015a.

Frostig, Roy, Ge, Rong, Kakade, Sham, and Sidford, Aaron. Competing with the empirical risk minimizer in a single pass. In *COLT*, 2015b.

Gürbüzbalaban, Mert, Ozdaglar, Asu, and Parrilo, Pablo. Why random reshuffling beats stochastic gradient descent. *arXiv preprint arXiv:1510.08560*, 2015.

Hofmann, Thomas, Lucchi, Aurelien, and McWilliams, Brian. Neighborhood watch: Stochastic gradient descent with neighbors. *NIPS*, 2015.

Hsieh, Cho-Jui, Chang, Kai-Wei, Lin, Chih-Jen, Keerthi, S. Sathiya, and Sundararajan, S. A dual coordinate descent method for large-scale linear SVM. In *ICML*, pp. 408–415, 2008.

Johnson, Rie and Zhang, Tong. Accelerating stochastic gradient descent using predictive variance reduction. In *NIPS*, pp. 315–323, 2013.

Konečnỳ, Jakub and Richtárik, Peter. Semi-stochastic gradient descent methods. *arXiv preprint arXiv:1312.1666*, 2013.

Lan, Guanghui. An optimal randomized incremental gradient method. *arXiv preprint arXiv:1507.02000*, 2015.

Lin, Hongzhou, Mairal, Julien, and Harchaoui, Zaid. A universal catalyst for first-order optimization. In *NIPS*, pp. 3366–3374, 2015a.

Lin, Qihang, Lu, Zhaosong, and Xiao, Lin. An accelerated randomized proximal coordinate gradient method and its application to regularized empirical risk minimization. *SIAM Journal on Optimization*, 25(4):2244–2273, 2015b.

Mairal, Julien. Incremental majorization-minimization optimization with application to large-scale machine learning. *SIAM Journal on Optimization*, 25(2):829–855, 2015.

Nitanda, Atsushi. Stochastic proximal gradient descent with acceleration techniques. In *NIPS*, pp. 1574–1582, 2014.

Pollard, David. *Convergence of stochastic processes*. Springer-Verlag, 1984.

Qu, Zheng, Richtarik, Peter, and Zhang, Tong. Randomized dual coordinate ascent with arbitrary sampling. *arXiv:1411.5873*, 2014.

Rakhlin, Alexander, Shamir, Ohad, and Sridharan, Karthik. Making gradient descent optimal for strongly convex stochastic optimization. In *ICML*, 2012.

Recht, Benjamin and Ré, Christopher. Beneath the valley of the noncommutative arithmetic-geometric mean inequality: conjectures, case-studies, and consequences. *COLT*, 2012.

Roux, Nicolas Le, Schmidt, Mark W., and Bach, Francis. A stochastic gradient method with an exponential convergence rate for finite training sets. In *NIPS*, pp. 2672–2680, 2012.

Schmidt, Mark, Roux, Nicolas Le, and Bach, Francis. Minimizing finite sums with the stochastic average gradient, 2013.

Shalev-Shwartz, Shai and Srebro, Nathan. Svm optimization: inverse dependence on training set size. In *ICML*, pp. 928–935, 2008.

Shalev-Shwartz, Shai and Zhang, Tong. Stochastic dual coordinate ascent methods for regularized loss. *Journal of Machine Learning Research*, 14(1):567–599, 2013.

Shalev-Shwartz, Shai and Zhang, Tong. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. In *ICML*, pp. 64–72, 2014.

Shalev-Shwartz, Shai, Singer, Yoram, Srebro, Nathan, and Cotter, Andrew. Pegasos: primal estimated sub-gradient solver for SVM. *Math. Program.*, 127(1):3–30, 2011.

Sridharan, Karthik, Srebro, Nathan, and Shalev-Shwartz, Shai. Fast rates for regularized objectives. In *NIPS*, pp. 1545–1552, 2009.

Suzuki, Taiji. Stochastic dual coordinate ascent with alternating direction method of multipliers. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pp. 736–744, 2014.

Vainsencher, Daniel, Liu, Han, and Zhang, Tong. Local smoothness in variance reduced optimization. In *NIPS*, pp. 2170–2178, 2015.

Xiao, Lin and Zhang, Tong. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.

Zhang, Lijun, Mahdavi, Mehrdad, and Jin, Rong. Linear convergence with condition number independent access of full gradients. In *Advances in Neural Information Processing Systems*, pp. 980–988, 2013.

Zhang, Yuchen and Xiao, Lin. Stochastic primal-dual coordinate method for regularized empirical risk minimization. In *ICML*, 2015.

Zhao, Peilin and Zhang, Tong. Stochastic optimization with importance sampling. *arXiv:1401.2753*, 2014.

Zhu, Zeyuan Allen, Qu, Zheng, Richtarik, Peter, and Yuan, Yang. Even faster accelerated coordinate descent using non-uniform sampling. *arXiv preprint arXiv:1512.09103*, 2015.

Appendix: Additional Empirical Results

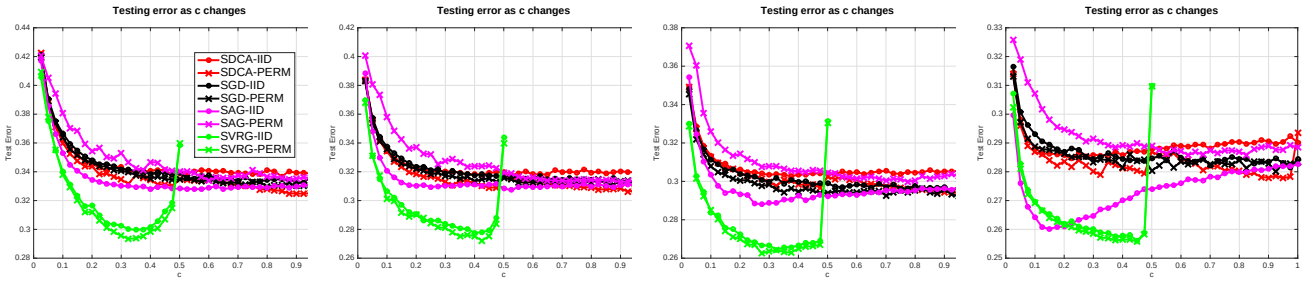

covtype, from left to right: T = 1000,2000,4000,8000

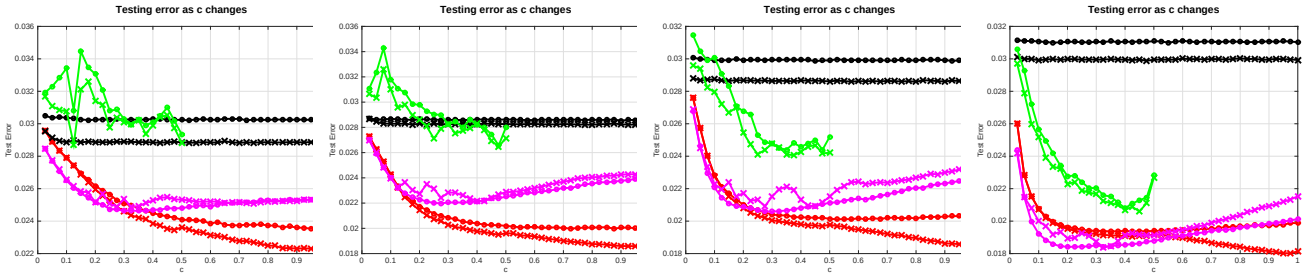

w8a, from left to right: T = 4000,8000,16000,32000

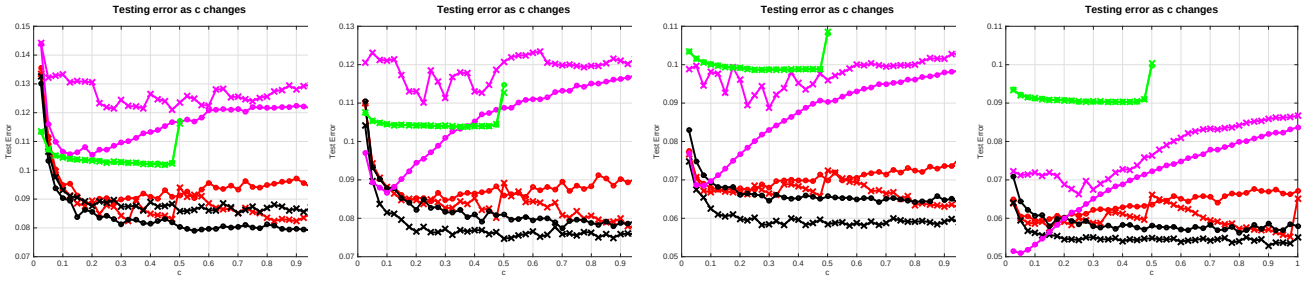

svmguide1, from left to right: T = 500,1000,2000,4000

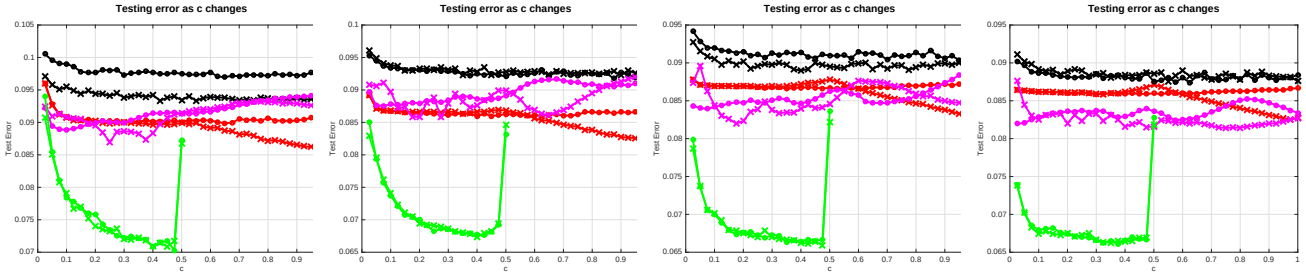

ijcnn1, from left to right: T = 4000,8000,16000,32000

Figure 7. Illustration of generalization errors as c varied